

Locality and Availability of Array Codes Constructed from Subspaces

Natalia Silberstein^{*†}, Tuvi Etzion[†], and Moshe Schwartz^{*}

^{*}Electrical and Computer Engineering, Ben-Gurion University of the Negev,

Beer Sheva 8410501, Israel, schwartz@ee.bgu.ac.il

[†]Computer Science, Technion – Israel Institute of Technology,

Haifa 3200003, Israel, etzion@cs.technion.ac.il

[‡]Yahoo! Labs, Haifa 31905, Israel, natalys@cs.technion.ac.il

Abstract—Ever-increasing amounts of data are created and processed in internet-scale companies such as Google, Facebook, and Amazon. The efficient storage of such copious amounts of data has thus become a fundamental and acute problem in modern computing. No single machine can possibly satisfy such immense storage demands. Therefore, distributed storage systems (DSS), which rely on tens of thousands of storage nodes, are the only viable solution. Such systems are broadly used in all modern internet-scale systems. However, the design of a DSS poses a number of crucial challenges, markedly different from single-user storage systems. Such systems must be able to reconstruct the data efficiently, to overcome failure of servers, to correct errors, etc. Lots of research was done in the last few years to answer these challenges and the research is increasing in parallel to the increasing amount of stored data.

The main goal of this paper is to consider codes which have two of the most important features of distributed storage systems, namely, locality and availability. Our codes are array codes which are based on subspaces of a linear space over a finite field. We present several constructions of such codes which are q -analog to some of the known block codes. Some of these codes possess independent intellectual merit. We examine the locality and availability of the constructed codes. In particular we distinguish between two types of locality and availability, node vs. symbol, locality and availability. To our knowledge this is the first time that such a distinction is given in the literature.

I. INTRODUCTION

Designing efficient mechanisms to store, maintain, and efficiently access large volumes of data is a highly relevant problem. Indeed, ever-increasing amounts of information are being generated and processed in the data centers of Amazon, Facebook, Google, Dropbox, and many others. The demand for ever-increasing amounts of cloud storage is supplied through the use of Distributed Storage Systems (DSS), where data is stored on a network of nodes (hard drives and solid-state drives).

In the DSS paradigm, it is essential to store data redundantly, in order to tolerate inevitable node failures [1], [11], [23]. Currently, the resilience against node failures is typically afforded by *replication*, where several copies of each data object are stored on different storage nodes.

However, replication is highly inefficient in terms of storage capacity. Recently, *erasure-correcting codes* have been used in DSS to reduce the large storage overhead of replicated systems [3], [5], [15].

Apart from storage space, other metrics should be considered when designing an actual DSS. However, in contrast with storage space, these metrics are adversely affected by the straightforward use of simple erasure-correcting codes. One such metric is the *repair bandwidth*: the amount of data that needs to be transferred when a node has failed, and is thus replaced. This metric is highly relevant as a prohibitively large fraction of the network bandwidth in a DSS may be consumed by such repair operations. Let us term *all* the information stored by a DSS as *the file*. Traditional erasure-correcting codes, and in particular *maximum distance separable (MDS)* codes, usually require that *all* the file be downloaded in order to regenerate a failed node. Recently, Dimakis et al. [4] established a tradeoff between the repair bandwidth and the storage capacity of a node, and introduced a new family of erasure-correcting codes, called *regenerating codes*, which attain this tradeoff. In particular, they proved that if a *large number* of storage nodes can be contacted during the repair of a failed node, and only a *fraction of their stored data* is downloaded, then the repair bandwidth can be minimized.

Local repairability of a DSS is an additional property which is highly sought. The corresponding performance metric is termed the *locality* of the coding scheme: the number of nodes that must participate in a repair process when a particular node fails. Local repairability is of significant interest when a cost is associated with contacting each node in the system. This is indeed the case in real world scenarios, for example as the result of network constraints. Codes which enable local repairs of failed system nodes are called *locally repairable codes (LRCs)*. These codes were introduced by Gopalan et al. in [12]. LRCs which also minimize the repair bandwidth, called codes with local regeneration, were considered in [16], [17], [21].

Regenerating codes and LRCs are attractive primarily for the storage of *cold* data: archival data that is rarely accessed. On the other hand, they do not address the challenges posed by the storage of frequently accessed *hot* data. For example, hot-data storage must enable efficient reads

This work was supported in part by the Israeli Science Foundation (ISF), Jerusalem, Israel, under Grant no. 10/12, and by the Israeli Science Foundation (ISF), Jerusalem, Israel, under Grant no. 130/14.

of the same data segments by several users *in parallel*. This property is referred to as *availability*. Codes which provide both locality and availability were first proposed in [22].

Regenerating codes are described in terms of stored information in nodes (servers). In other words, regenerating codes are usually array codes [25]. Reconstructing the files and repairing failed nodes are the main tasks of regenerating codes. LRCs and codes with availability are usually described as block codes, and access and/or repair is described in terms of symbols.

In this work we combine the two approaches and discuss two types of locality (availability, respectively), node locality (availability) which resembles the first approach and symbol locality (availability) which resembles the second approach. To our knowledge, such a combined approach was not considered in the literature before.

Our solution approach will be based on array codes, constructed via subspaces of a finite vector space. A subspace approach for DSS was considered for the first time in [13] and later in [20]. Our approach is slightly different from the approach in these two papers. We will design array codes based on subspaces and analyze their locality and availability properties.

The rest of this paper is organized as follows. Preliminaries are given in Section II. Our subspace approach, constructions of codes, and analysis of their locality and availability, are presented in Section III. We first describe a generalization of the simplex code in Section III-A, and then extend this approach with further codes and their duals in Section III-B.

II. PRELIMINARIES

Let \mathbb{F}_q denote the finite field of size q . For a natural number $m \in \mathbb{N}$, we use the notation $[m] \triangleq \{1, 2, \dots, m\}$. We use lower-case letters to denote scalars. Overlined letters denote vectors, which by default are assumed to be column vectors. Matrices are denoted by upper-case letters. Most literature denotes codewords (which are usually vectors) by overlined lower-case letters. However, since we also have codewords which are arrays (matrices), these will be denoted by bold lower-case letters. Thus, typically, we shall have a generator matrix G , whose j th column is \bar{g}_j , and whose (i, j) th entry is $g_{i,j}$. An array code will usually be denoted by C , whose typical codeword will be denoted by \mathbf{c} . We use 0 to denote the scalar zero, $\bar{0}$ for the all-zero vector, and $\mathbf{0}$ for the all-zero matrix. Also, given a (possibly empty) set of vectors, $\bar{v}_1, \dots, \bar{v}_m \in \mathbb{F}_q^n$, their span is denoted by $\langle \bar{v}_1, \dots, \bar{v}_m \rangle$.

Our main object of study is a linear array code, formally defined as follows: A $[b \times n, M, d]$ array code over \mathbb{F}_q , denoted C , is a linear subspace of $b \times n$ matrices over \mathbb{F}_q . Matrices $\mathbf{c} \in C$ are referred to as *codewords*. The elements of a codeword are denoted by $c_{i,j}$, $i \in [b]$, $j \in [n]$, and are referred to as *symbols*. Columns of codewords are denoted by \bar{c}_j , $j \in [n]$. We denote by $M \triangleq \dim(C)$ the *dimension* of the code as a linear space over \mathbb{F}_q . The *weight* of an

array is defined as the number of non-zero columns, i.e., for $\mathbf{c} \in C$,

$$\text{wt}(\mathbf{c}) \triangleq |\{\bar{c}_j : \bar{c}_j \neq \bar{0}, j \in [n]\}|.$$

Finally, the *minimum distance* of the code, denoted d , is the defined as the minimal weight of a non-zero codeword,

$$d \triangleq \min_{\substack{\mathbf{c} \in C \\ \mathbf{c} \neq \mathbf{0}}} \text{wt}(\mathbf{c}).$$

We make two observations to avoid confusion with other notions of error-correcting codes. The first observation is that by reading the symbols of codewords, column by column, and within each column, from first to last entry, we may flatten the $b \times n$ codewords to vectors of length bn . This results in a code over \mathbb{F}_q of length bn , dimension M , but more often than not, a larger minimum distance, since the above definition considers non-zero *columns* and not non-zero symbols. Assume G is an $M \times bn$ generator matrix for the flattened code. By abuse of notation, we shall also call G the *generator matrix* for the original code C . Note that in G , columns $(j-1)b+1, \dots, jb$, correspond to the symbols appearing in the j th codeword column in C . We shall call these b columns in G by the j th *thick column* of G , similarly to [16]. Thus, G is a matrix comprised of n thick columns, corresponding to the n columns of codewords in C .

The second observation is that we may use the well known isomorphism $\mathbb{F}_q^b \cong \mathbb{F}_{q^b}$, and consider each column of a codeword as a single element from \mathbb{F}_{q^b} . We get an \mathbb{F}_q -linear code over \mathbb{F}_{q^b} (sometimes called a *vector-linear code*), of length n , minimum distance d , but with a dimension (taken as usual over \mathbb{F}_{q^b}) not necessarily M .

In a typical distributed-storage setup, we would like to store a file containing M symbols from \mathbb{F}_q . The file is encoded into an array $\mathbf{c} \in C$ from a $[b \times n, M, d]$ array code. Each codeword column of \mathbf{c} is stored in a different node. The minimum distance d of the code ensures that any failure of at most $d-1$ nodes may be corrected.

Two important properties of codes for distributed storage are *locality* and *availability*. An important feature of this paper is the distinction between *symbol* locality and *node* locality (respectively, availability).

Let C be a $[b \times n, M, d]$ array code. We say a codeword column $j \in [n]$ has *node locality* r_n , if its content may be obtained via linear combinations of the contents of the recovery-set columns. More precisely, there exists a recovery set $S = \{j_1, \dots, j_{r_n}\} \subseteq [n] \setminus \{j\}$ of r_n other codeword columns, and scalars $a_{\ell,m}^{(i)} \in \mathbb{F}_q$, $i, \ell \in [b]$, $m \in [r_n]$, such that for all $i \in [b]$,

$$c_{i,j} = \sum_{\ell=1}^b \sum_{m=1}^{r_n} a_{\ell,m}^{(i)} c_{\ell,j_m} \quad (1)$$

simultaneously for all codewords $\mathbf{c} \in C$. If all codeword columns have this property, we say the code has node locality of r_n . We emphasize the fact the a column is not necessarily a linear combination of its recovery-set columns. Rather, each of its symbols may be obtained by

a linear combination of the symbols of its recovery-set columns (possibly a different linear combination for each symbol).

Similarly, we say the code has *symbol locality* r_s , if for every coordinate, $i \in [b]$ and $j \in [n]$, there exists a recovery set $S = \{j_1, \dots, j_{r_s}\} \subseteq [n] \setminus \{j\}$ of r_s other codeword columns, and scalars $a_{\ell,m} \in \mathbb{F}_q$, $\ell \in [b]$, $m \in [r_s]$, such that for every codeword $\mathbf{c} \in C$,

$$c_{i,j} = \sum_{m=1}^{r_s} \sum_{\ell=1}^b a_{\ell,m} c_{\ell,j_m}. \quad (2)$$

Thus, each code symbol may be recovered from the code symbols in r_s other codeword columns. It is obvious that $r_s \leq r_n$.

Once locality is defined, we can also define availability. The *node availability*, denoted t_n , (respectively, the *symbol availability*, denoted t_s) is the number of pairwise-disjoint recovery sets (as in the definition of locality) that exist for any codeword column (respectively, symbol). Note that each recovery set should of size at most r_n (respectively, r_s).

We also recall some useful facts regarding Gaussian coefficients. Let V be a vector space of dimension n over \mathbb{F}_q . For any integer $0 \leq k \leq n$, we denote by $\begin{bmatrix} V \\ k \end{bmatrix}$ the set of all k -dimensional subspaces of V . The *Gaussian coefficient* is defined for n , k , and q as

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \frac{(q^n - 1)(q^{n-1} - 1) \dots (q^{n-k+1} - 1)}{(q^k - 1)(q^{k-1} - 1) \dots (q - 1)}.$$

Whenever the size of the field, q , is clear from the context, we shall remove the subscript q .

It is well known that the number of k -dimensional subspaces of an n -dimensional space over \mathbb{F}_q is given by $\begin{bmatrix} n \\ k \end{bmatrix}$. The Gaussian coefficients satisfy the following recursions (see [26, Chapter 24]),

$$\begin{aligned} \begin{bmatrix} n \\ k \end{bmatrix} &= \begin{bmatrix} n-1 \\ k \end{bmatrix} + q^{n-k} \begin{bmatrix} n-1 \\ k-1 \end{bmatrix} \\ &= q^k \begin{bmatrix} n-1 \\ k \end{bmatrix} + \begin{bmatrix} n-1 \\ k-1 \end{bmatrix}. \end{aligned} \quad (3)$$

Additionally, in a more general form, the number of k' -dimensional subspaces of V which intersect a given k -dimensional subspace of V in an i -dimensional subspace is given by (see [9])

$$q^{(k'-i)(k-i)} \begin{bmatrix} n-k \\ k'-i \end{bmatrix} \begin{bmatrix} k \\ i \end{bmatrix}. \quad (4)$$

III. A SUBSPACE APPROACH TO LRCs

Let C be a $[b \times n, M, d]$ array code over \mathbb{F}_q . We now describe an approach to viewing such array codes which will lead to the main results of this section.

Denote $V \triangleq \mathbb{F}_q^M$ the M -dimensional vector space over \mathbb{F}_q . Let G be a generator matrix for the (flattened) array code C . For each $j \in [n]$, we define $V_j \in \bigcup_{k=0}^b \begin{bmatrix} V \\ k \end{bmatrix}$, to be the column space of the j th thick column of G , i.e.,

$$V_j \triangleq \langle \bar{g}_{(j-1)b+1}, \bar{g}_{(j-1)b+2}, \dots, \bar{g}_{jb} \rangle.$$

We say V_j is associated with the j th thick column of G , or equivalently, associated with the j th column of the codewords of C .

The following equivalence is fundamental to the constructions and analysis of this section.

Lemma 1. *Let C be a $[b \times n, M, d]$ array code over \mathbb{F}_q , and let V_j , $j \in [n]$, be the subspaces associated with the codeword columns. Then $S = \{j_1, \dots, j_m\} \subseteq [n] \setminus \{j\}$ is a recovery set for codeword column $j \in [n]$, if and only if*

$$V_j \subseteq V_{j_1} + V_{j_2} + \dots + V_{j_m}.$$

Similarly, S is a recovery set for symbol (i, j) , $i \in [b]$, if

$$\bar{g}_{(j-1)b+i} \in V_{j_1} + V_{j_2} + \dots + V_{j_m},$$

where $\bar{g}_{(j-1)b+i}$ is the i th column in the j th thick column of a generating matrix G for C .

Proof: This is a simple restatement of (1) and (2). ■

With this equivalence, we may obtain the node/symbol locality/availability using subspace properties of the thick columns of a generating matrix. Another definition of interest is the following.

Definition 2. *Let C be a $[b \times n, M, d]$ array code over \mathbb{F}_q , and let V_j be the subspace associated with the j th thick column. If $\dim(V_j) = b$ for all $j \in [n]$ we call C full column rank.*

A. Generalized Simplex Codes via Subspaces

We start with a construction of array codes which may be considered as a generalization and a q -analog of the classical simplex code, the dual of the Hamming code (see [18, p. 30]).

Construction A. *Fix a finite field \mathbb{F}_q , positive integers $b \leq M$, and $V = \mathbb{F}_q^M$. Construct a $b \times \begin{bmatrix} M \\ b \end{bmatrix}$ array code whose set of columns are associated with the subspaces $\begin{bmatrix} V \\ b \end{bmatrix}$, each appearing exactly once.*

We make a note here, which is also relevant for the constructions to follow. Once we fix the set of subspaces associated with the codeword columns, the code is constructed in the following way: for each $j \in [n]$, and associated subspace V_j , we arbitrarily choose a set of b vectors from \mathbb{F}_q^M that form a basis for V_j . These b vectors are placed (in some arbitrary order) as the columns comprising the j th thick column of a generator matrix G . The resulting matrix G generates the constructed code¹.

Lemma 3. *Fix a finite field \mathbb{F}_q , positive integers $b < M$, and $V = \mathbb{F}_q^{M-1}$. For any $V' \in \begin{bmatrix} V \\ b-1 \end{bmatrix}$, given as the column space of an $(M-1) \times (b-1)$ matrix G' , and for any non-zero vector $\bar{u} \in \mathbb{F}_q^{M-1}$ such that $\bar{u}^T G' = \bar{0}^T$, the following hold:*

- 1) *If $\bar{x}, \bar{y} \in \mathbb{F}_q^{M-1}$ are in the same coset of V' , then $\bar{u}^T \bar{x} = \bar{u}^T \bar{y}$.*

¹The choice of basis vectors for each V_j , their order, and the order of thick columns, is arbitrary. Distinct choices result in equivalent codes.

- 2) The number of cosets of V' , all of whose vectors \bar{x} satisfy $\bar{u}^T \bar{x} = a$, for some fixed $a \in \mathbb{F}_q$, is exactly q^{M-b-1} .

Proof: Denote the columns of G' as $\bar{g}'_1, \dots, \bar{g}'_{b-1}$. If \bar{x} and \bar{y} are in the same coset of V' , then there exist scalars a_1, \dots, a_{b-1} such that

$$\bar{x} = \bar{y} + \sum_{j=1}^{b-1} a_j \bar{g}'_j.$$

Multiplying on the left by \bar{u}^T , and recalling that $\bar{u}^T G' = \bar{0}^T$, we obtain the first claim.

The number of cosets of V' is exactly q^{M-b} , each containing q^{b-1} vectors. Since $\bar{u} \neq \bar{0}$, the number of vectors $\bar{x} \in \mathbb{F}_q^{M-1}$ such that $\bar{u}^T \bar{x} = a$ is q^{M-2} . Dividing this by the number of vectors per coset we obtain the second claim. ■

We are now ready for the first claim on the properties of the codes from Construction A.

Theorem 4. The array code obtained from Construction A is a $[b \times \binom{M}{b}, M, d]$ array code, with

$$d = \binom{M}{b} - \binom{M-1}{b} = q^{M-b} \binom{M-1}{b-1}.$$

Additionally, except for the all-zero array codeword, all other codewords have the same constant weight d .

Proof: Apart from the minimum distance of the code, all other parameters are trivial. We shall prove the minimum distance property by proving the constant-weight property of the non-zero codewords by induction on M and b . To make the dependence on the code parameters explicit, we denote the code from Construction A by C_b^M . Additionally, we assert an auxiliary claim on the thick columns of the generator matrix, namely, that each thick column has rank b . We will prove this claim by induction as well.

For the induction basis, we have the following cases. When considering C_M^M , the codewords are $M \times 1$ arrays, and trivially, any non-zero codeword has weight

$$1 = q^{M-M} \binom{M-1}{M-1}.$$

Another base case is C_1^M . In the resulting generator matrix, each thick column contains just a single column, and the matrix is nothing but a generator matrix for the well known simplex code. The codewords are $1 \times (q^M - 1)/(q - 1)$ arrays. The weight of the non-zero codewords in the simplex code is known to be q^{M-1} , and indeed we get a constant weight of

$$q^{M-1} = q^{M-1} \binom{M-1}{0}.$$

We additionally note that in both cases, each thick column has rank b .

Assume now the claim holds for C_{b-1}^{M-1} and for C_b^{M-1} . For the induction step we prove the claim also holds for C_b^M . Let their respective generating matrices be G_{b-1}^{M-1} and

G_b^{M-1} . Since we are not in any of the induction-base cases, we additionally have $1 < b < M$.

We construct a new matrix, G by concatenating modified thick columns from G_{b-1}^{M-1} and G_b^{M-1} . We first take each thick column of G_b^{M-1} , append a bottom row of all zeros, and place it as a thick column of G . We call these columns *thick columns of type I*.

All the remaining thick columns of G , which we call of *type II*, are formed by the thick columns of G_{b-1}^{M-1} as follows. Consider such a single thick column, which is an $(M-1) \times (b-1)$ matrix on its own. Denote its column space by $V' \subseteq \mathbb{F}_q^{M-1}$, which by the induction hypothesis, has rank $b-1$. Thus, there are q^{M-b} cosets of V' in \mathbb{F}_q^{M-1} . Let $\bar{v}'_1, \dots, \bar{v}'_{q^{M-b}}$ be arbitrary coset representatives of the distinct cosets of V' . We create q^{M-b} thick columns in G from the given thick column of G_{b-1}^{M-1} by placing it, each time with \bar{v}'_i as a b th column, and with an appended bottom row of $0, \dots, 0, 1$. In such thick columns of type II, the left $b-1$ coordinates are called *the recursive part*, whereas the last coordinate is called *the coset part*. The two types of thick columns of G (depending on their source) are depicted in Figure 1.




Figure 1. The two types of thick columns in the constructed matrix G : a type I thick column, created by a thick column (a) from G_b^{M-1} , and a type II thick column, created by a thick column (b) from G_{b-1}^{M-1} and one of its column-space coset representatives.

Simple bookkeeping shows that we have $\binom{M-1}{b}$ columns of type I, and $q^{M-b} \binom{M-1}{b-1}$ columns of type II, for a total of

$$\binom{M-1}{b} + q^{M-b} \binom{M-1}{b-1} = \binom{M}{b}$$

columns, where we used (3). They are easily seen to have distinct associated subspaces, each of dimension b , accounting for all the b -dimensional subspaces of $V = \mathbb{F}_q^M$. Thus, G is indeed a generator matrix for the code from Construction A, where each column has rank b .

Now that we have proven a decomposition for the generator matrix G , we can proceed with the proof of the constant weight of all non-zero codewords. It is easily seen that G has full rank. We consider several cases, depending on the rows of G participating in the linear combination creating the codeword at question.

In the simplest case, if a codeword of C_b^M is formed by the last row of G only, then its weight is $q^{M-b} \binom{M-1}{b-1}$, as the number of thick columns of type II.

For the second case, let us consider a codeword $\mathbf{c} \in C_b^M$ formed by a linear combination of some rows from the first $M-1$ rows of G . By the induction hypothesis, the thick columns of type I contribute $\binom{M-1}{b} - \binom{M-2}{b}$ to the weight of \mathbf{c} . Also by the induction hypothesis, the recursive parts of

thick columns of type II contribute $q^{M-b}(\binom{M-1}{b-1} - \binom{M-2}{b-1})$ to the weight. Finally, even if for some thick column of type II the recursive part may produce a combination of all zeros, the coset part may be non-zero, thus contributing to the weight of \mathbf{c} . More precisely, we have $\binom{M-2}{b-1}$ recursive parts the linear combination zeros. Therefore, by Lemma 3, the coset part of exactly $\binom{M-2}{b-1}(q-1)q^{M-b-1}$ becomes non-zero, and contributes to the weight of \mathbf{c} . In total we get,

$$\begin{aligned} \text{wt}(\mathbf{c}) &= \binom{M-1}{b} - \binom{M-2}{b} \\ &\quad + q^{M-b} \left(\binom{M-1}{b-1} - \binom{M-2}{b-1} \right) \\ &\quad + \binom{M-2}{b-1} (q-1)q^{M-b-1} \\ &= \binom{M}{b} - \binom{M-1}{b}. \end{aligned}$$

Finally, we consider a linear combination that, non-trivially, uses some rows from the set of $M-1$ first rows, as well as the last row. The 1's in the last row are located exactly at the coset part of thick columns of type II. Since by Lemma 3, the linear combination results in an equal number of appearances of each element of \mathbb{F}_q in the coset parts, an addition of a multiple of the last row will not change that, and the weight of the codeword remains the same as in the previous case. ■

We observe that the codes of Construction A form a generalization of simplex codes. When we choose $b = 1$ in Construction A, the simplex code is obtained, a fact that was used in the proof of Theorem 4.

Lemma 5. *The array code obtained from Construction A, with parameters $b < M$, has node locality of $r_n = 2$, and symbol locality of*

$$r_s = \begin{cases} 1 & b > 1, \\ 2 & b = 1. \end{cases}$$

Proof: Let C be a code generated by Construction A with a generator matrix G . We first examine the case of $b > 1$. For symbol locality, given any column of G , denoted $\bar{g} \in \mathbb{F}_q^M$, by (4), there are exactly $\binom{M-1}{b-1}$ b -dimensional subspaces of \mathbb{F}_q^M containing \bar{g} , each corresponding to a thick column of G . Since $b < M$, we have $\binom{M-1}{b-1} > 1$, and there exists a thick column different than the one containing the column \bar{g} , whose column space contains \bar{g} . Hence, $r_s = 1$.

For node locality, given any subspace V_j associated with the j th thick column of G , we can easily find two other subspaces V_{j_1} and V_{j_2} , $j \notin \{j_1, j_2\}$, such that $V_j \subseteq V_{j_1} + V_{j_2}$. For example: fix a basis for V_j . Take the first basis element and complete it to a basis of some b -dimensional subspace of \mathbb{F}_q^M , denoted V_{j_1} . Take the remaining $b-1$ basis elements of V_j and complete them to a different b -dimensional subspace, denoted V_{j_2} . This can always be done when $1 < b < M$. Hence, $r_n = 2$.

Finally, we consider the case $b = 1$. In this case, each thick column of G comprises of a single column. By

definition this means that $r_n = r_s$, and since each column may be shown as the sum of two other columns, we have $r_n = r_s = 2$. ■

We note that we ignored the case of $b = M$ in the previous lemma, since then the array codewords have a single column, and locality is not defined.

We now turn to consider availability. Symbol availability is trivial.

Corollary 6. *The array code obtained from Construction A, with parameters $1 < b < M$, has symbol availability*

$$t_s = \binom{M-1}{b-1} - 1$$

and for $b = 1$ $t_s = \frac{q^{M-1}-1}{2}$.

Proof: We use (4) to find the number of associated subspaces containing a given vector. ■

Unlike locality, it appears that determining the node availability is a difficult task. We consider only the simplest non-trivial case of $b = 2$.

Lemma 7. *The array code obtained from Construction A, with parameters $b = 2 < M$, has node availability*

$$t_n = \frac{1}{2} \left(\binom{M}{2} - 1 \right),$$

when q is even, and

$$t_n \geq \frac{1}{2} \left(\binom{M}{2} - 1 - q(q^2 + q - 1) \binom{M-2}{2} \right),$$

when q is odd.

Proof: Let us consider some codeword column of the code, and its associated subspace, $V = \langle \bar{v}_1, \bar{v}_2 \rangle$. We count the number of pairwise-disjoint pairs of subspaces $U, W \neq V$, such that $V \subseteq U + W$. We show how all subspaces (except for V) may be paired in such a manner, except perhaps for a few due to parity issues. We distinguish between two different kinds of subspaces, where the subspaces of the first kind intersect V in a one-dimensional subspace (a projective point), and where the subspaces of the second kind have only trivial intersection with V .

First, we consider subspaces of the first kind. There are $\binom{M-1}{1} - 1 = q \binom{M-2}{1}$ associated subspaces different from V that contain a given vector $\bar{v} \in V$, $\bar{v} \neq \bar{0}$, and we denote them by $\mathcal{V}_{\bar{v}}$. Since there are $\binom{2}{1} = q+1$ projective points in V , denoted $\bar{v}_1, \dots, \bar{v}_{q+1}$, we have $q(q+1) \binom{M-2}{1}$ associated subspaces which intersect V in a one-dimensional subspace. Note that if $U \in \mathcal{V}_{\bar{v}_i}$ and $W \in \mathcal{V}_{\bar{v}_j}$, with $i \neq j$, then $V \subseteq U + W$. We now further partition each $\mathcal{V}_{\bar{v}_i}$ into q sets of equal size, arbitrarily. We denote these $\mathcal{V}_{\bar{v}_i}^j$, where $j \in [q+1] \setminus \{i\}$. The size of each such set is

$$|\mathcal{V}_{\bar{v}_i}^j| = \binom{M-2}{1}.$$

Finally, for each $i, j \in [q+1]$, $i \neq j$, we arbitrarily create pairs of elements, one from $\mathcal{V}_{\bar{v}_i}^j$, and one from $\mathcal{V}_{\bar{v}_j}^i$. The total number of such pairs is $\binom{q+1}{2} \binom{M-2}{1}$.

Next we consider associated subspaces of the second kind. There are $\binom{M}{2} - 1 - q(q+1)\binom{M-2}{1}$ such subspaces. We will prove that for even q one can partition all these subspaces into disjoint pairs, and for odd q one can partition all but a few such subspaces into disjoint pairs. The statement of the lemma then follows from this proof.

Given an associated subspace $U = \langle \bar{u}_1, \bar{u}_2 \rangle$, $U \cap V = \{\bar{0}\}$, we define a set \mathcal{S}_U of q^4 subspaces, as follows:

$$\mathcal{S}_U = \{ \langle \bar{u}_1 + \bar{x}_1, \bar{u}_2 + \bar{x}_2 \rangle : \bar{x}_1, \bar{x}_2 \in V \}.$$

Note that since $U \cap V = \{\bar{0}\}$, the vectors $\bar{u}_1 + \bar{x}_1$ and $\bar{u}_2 + \bar{x}_2$ are linearly independent. One can easily verify that \mathcal{S}_U is well defined, and the choice of two basis vectors, \bar{u}_1 and \bar{u}_2 , does not change \mathcal{S}_U .

Additionally, if we have two distinct associated subspaces of the second kind, $U \neq U'$, then either $\mathcal{S}_U \cap \mathcal{S}_{U'} = \emptyset$ or $\mathcal{S}_U = \mathcal{S}_{U'}$. To see that, assume $W_1 \in \mathcal{S}_U \cap \mathcal{S}_{U'}$, i.e.,

$$\begin{aligned} W_1 &= \langle \bar{u}_1 + \bar{x}_1, \bar{u}_2 + \bar{x}_2 \rangle \in \mathcal{S}_U, \\ W_1 &= \langle \bar{u}'_1 + \bar{x}'_1, \bar{u}'_2 + \bar{x}'_2 \rangle \in \mathcal{S}_{U'}, \end{aligned}$$

with $\bar{x}_1, \bar{x}_2, \bar{x}'_1, \bar{x}'_2 \in V$. Then there exist $\alpha_{1,1}, \alpha_{1,2}, \alpha_{2,1}, \alpha_{2,2} \in \mathbb{F}_q$ such that

$$\begin{aligned} \bar{u}_1 + \bar{x}_1 &= \alpha_{1,1}(\bar{u}'_1 + \bar{x}'_1) + \alpha_{1,2}(\bar{u}'_2 + \bar{x}'_2), \\ \bar{u}_2 + \bar{x}_2 &= \alpha_{2,1}(\bar{u}'_1 + \bar{x}'_1) + \alpha_{2,2}(\bar{u}'_2 + \bar{x}'_2), \end{aligned}$$

and

$$\Delta = \det \begin{pmatrix} \alpha_{1,1} & \alpha_{1,2} \\ \alpha_{2,1} & \alpha_{2,2} \end{pmatrix} \neq 0.$$

We cannot have $\alpha_{1,1} = \alpha_{1,2} = 0$, and we assume $\alpha_{1,2} \neq 0$ where the other case is symmetric. Then, given $W_2 \in \mathcal{S}_U$, $W_2 = \langle \bar{u}_1 + \bar{y}_1, \bar{u}_2 + \bar{y}_2 \rangle$, where $\bar{y}_1, \bar{y}_2 \in V$, we define

$$\begin{aligned} \bar{y}'_1 &\triangleq \bar{x}'_1 + \frac{\alpha_{1,2}}{\Delta} \left(\frac{\alpha_{2,2}}{\alpha_{1,2}} (\bar{y}_1 - \bar{x}_1) - (\bar{y}_2 - \bar{x}_2) \right), \\ \bar{y}'_2 &\triangleq \bar{x}'_2 + \frac{1}{\alpha_{1,2}} (\bar{y}_1 - \bar{x}_1 - \alpha_{1,1}(\bar{y}'_1 - \bar{x}'_1)). \end{aligned}$$

Obviously, $\bar{y}'_1, \bar{y}'_2 \in V$. We also observe that

$$\begin{aligned} \bar{u}_1 + \bar{y}_1 &= \alpha_{1,1}(\bar{u}'_1 + \bar{y}'_1) + \alpha_{1,2}(\bar{u}'_2 + \bar{y}'_2), \\ \bar{u}_2 + \bar{y}_2 &= \alpha_{2,1}(\bar{u}'_1 + \bar{y}'_1) + \alpha_{2,2}(\bar{u}'_2 + \bar{y}'_2), \end{aligned}$$

and so $W_2 = \langle \bar{u}'_1 + \bar{y}'_1, \bar{u}'_2 + \bar{y}'_2 \rangle \in \mathcal{S}_{U'}$. Hence, if $\mathcal{S}_U \cap \mathcal{S}_{U'} \neq \emptyset$, then $\mathcal{S}_U = \mathcal{S}_{U'}$.

Thus, as U ranges over all associated subspaces of the second kind, \mathcal{S}_U partitions that set of subspaces into equivalence classes. We arbitrarily identify each such class with a subspace U , and a pair of basis vectors, $\bar{u}_1, \bar{u}_2 \in U$.

Depending on the parity of q we have two cases. First we consider even q . We partition each class \mathcal{S}_U , identified by U and $\bar{u}_1, \bar{u}_2 \in U$, into disjoint pairs as follows: We pair each

$$W = \langle \bar{u}_1 + \bar{x}_1, \bar{u}_2 + \bar{x}_2 \rangle \in \mathcal{S}_U,$$

with

$$f(W) = \langle \bar{u}_1 + \bar{x}_1 + \bar{v}_1, \bar{u}_2 + \bar{x}_2 + \bar{v}_2 \rangle \in \mathcal{S}_U.$$

Since q is even, this is indeed well defined since $f(f(W)) = W$. Additionally, the objective is met since

$$V = \langle \bar{v}_1, \bar{v}_2 \rangle \subseteq W + f(W).$$

When q is odd, we partition each class \mathcal{S}_U , identified by U and $\bar{u}_1, \bar{u}_2 \in U$, into disjoint pairs by pairing

$$W = \langle \bar{u}_1 + \bar{x}_1, \bar{u}_2 + \bar{x}_2 \rangle \in \mathcal{S}_U,$$

with

$$f(W) = \langle \bar{u}_1 - \bar{x}_1, \bar{u}_2 - \bar{x}_2 \rangle \in \mathcal{S}_U.$$

Except for $\bar{x}_1 = \bar{x}_2 = \bar{0}$, this is indeed a pairing since $f(f(W)) = W$. Additionally, whenever \bar{x}_1 and \bar{x}_2 are linearly independent, we have

$$V = \langle \bar{v}_1, \bar{v}_2 \rangle \subseteq W + f(W).$$

The number of such pairs is $\frac{1}{2}(q^2 - 1)(q^2 - q)$. Hence, we are not using $q(q^2 + q - 1)$ subspaces of the q^4 subspaces in \mathcal{S}_U , and there are $\binom{M-2}{2}$ sets \mathcal{S}_U . ■

B. Codes from Subspace Designs

In this subsection we focus on constructing codes by using certain subspace designs. We first present a different generalization of simplex codes by using spreads. The resulting code is known, and we analyze it for completeness, and for motivating another construction that uses subspace designs.

Consider a finite field \mathbb{F}_q and the vector space $V \triangleq \mathbb{F}_q^M$. A b -spread of V is a set $\{V_1, V_2, \dots, V_n\} \subseteq \binom{M}{b}$ such that $V_i \cap V_j = \{\bar{0}\}$ for all $i, j \in [n]$, $i \neq j$, and additionally, $\bigcup_{i \in [n]} V_i = V = \mathbb{F}_q^M$. Thus, except for the zero vector, $\bar{0}$, a spread is a partition of \mathbb{F}_q^M into subspaces. It is known that a b -spread exists if and only if $b|M$. Simple counting shows that the number of subspaces in a spread is

$$n = \frac{q^M - 1}{q^b - 1} = \frac{\binom{M}{1}}{\binom{b}{1}}.$$

Let us start with a code obtained from a single spread. This code was already described in [19], in the context of self-repairing codes, and we bring it here for completeness.

Construction B. Fix a finite field \mathbb{F}_q , positive integers $b|M$, and $V = \mathbb{F}_q^M$. Construct a $b \times \frac{\binom{M}{1}}{\binom{b}{1}}$ array code whose set of columns are associated with the subspaces of a b -spread of V , each appearing exactly once.

Theorem 8. The array code obtained from Construction B is a $[b \times \frac{\binom{M}{1}}{\binom{b}{1}}, M, q^{M-b}]$ array code. Additionally, except for the all-zero array codeword, all other codewords have the same constant weight.

Proof: Denote $u \triangleq \frac{\binom{M}{1}}{\binom{b}{1}}$. Consider an $M \times bu$ generator matrix G for the code C from Construction B. It contains u thick columns, each made up of b columns. Let G_i , $i \in [u]$, be the $M \times b$ submatrix of G containing the b columns of the i th thick column, i.e., $G = (G_1 | G_2 | \dots | G_u)$.

We now take each G_i , $i \in [u]$, and construct from it an $M \times (q^b - 1)$ matrix we call G_i^{ext} , whose columns are the

column space of G_i except for $\bar{0}$. We concatenate those to obtain the $M \times (q^M - 1)$ matrix

$$G^{\text{ext}} \triangleq (G_1^{\text{ext}} | G_2^{\text{ext}} | \dots | G_u^{\text{ext}}).$$

Since the thick columns of G form a b -spread of \mathbb{F}_q^M , the columns of G^{ext} contain each possible vector exactly once, except for $\bar{0}$.

We now observe that a row of G_i^{ext} is $\bar{0}^T$ iff it is $\bar{0}^T$ in G_i . Additionally, a non-zero row of G_i^{ext} contains exactly q^{b-1} occurrences of each non-zero element of \mathbb{F}_q . Finally, each non-zero element of \mathbb{F}_q appears q^{M-1} times in each row of G^{ext} . Thus, given a row of G^{ext} , exactly $q^{M-1}/q^{b-1} = q^{M-b}$ of its u thick columns are non-zero, implying the same for the corresponding row in G , and then the associated array codeword has weight q^{M-b} .

We now want to prove the same thing for every non-trivial linear combination of the rows of G . First, note that having a b -spread of \mathbb{F}_q^M is equivalent to having $\text{rank}(G_i) = b$, and $\text{rank}(G_i | G_j) = 2b$, for all $i, j \in [u]$, $i \neq j$. Consider a linear combination of rows i_1, i_2, \dots, i_ℓ of G , each with a non-zero coefficient, resulting in a row vector \bar{v}^T . Replace row i_ℓ of G by the vector \bar{v}^T to obtain a new matrix $G' = (G'_1 | G'_2 | \dots | G'_u)$. Since the rank is invariant to such operations, $\text{rank}(G'_i) = b$ and $\text{rank}(G'_i | G'_j) = 2b$ for all $i, j \in [u]$, $i \neq j$. Thus G' is equivalent to a b -spread (perhaps different from the original one induced by G). Using the same logic as before, exactly q^{M-b} of the thick columns of \bar{v}^T are non-zero, completing the proof. ■

Lemma 9. *The array code obtained from Construction B, $b < M$, has symbol locality $r_s = 2$, and its node locality satisfies $2 \leq r_n \leq b + 1$. Moreover, there exist such array codes with $r_n \leq M/b$.*

Proof: To prove the symbol locality, we note that any column of G can be presented as a linear combination of two other columns which belong to two other distinct thick columns. Otherwise, if these two columns belong to the same thick column, we obtain a contradiction to the definition of a spread. Thus, $r_s \leq 2$. We also obviously have $r_s \leq 2$, otherwise we contradict the partitioning property of the spread.

For the node locality, since in general $r_s \leq r_n$ we have that $2 \leq r_n$. Let $\{\bar{v}_1, \dots, \bar{v}_b\}$ be a basis for a thick column of G which represents an element (subspace) V_i of the spread. Take an arbitrary $\bar{w} \notin V_i$ and define $\bar{u}_i \triangleq \bar{v}_i + \bar{w}$, for all $i \in [b]$. Observe that \bar{w} and all the vectors \bar{u}_i , $i \in [b]$, belong to $b + 1$ different subspaces (corresponding to thick columns) in a spread, or else these would intersect V_i non-trivially. Clearly, V_i can be reconstructed from these $b + 1$ subspaces.

For the remainder of the proof let us assume that the spread is constructed in a specific way, inferred from [7], given in more detail in [10], and described as follows. Every element (subspace) in the constructed spread is presented as the row space of a row-reduced echelon-form $b \times M$ matrix $[0 | 0 | \dots | 0 | I_b | A_1 | A_2 | \dots | A_t]$, where each

block is of size $b \times b$, I_b is the $b \times b$, identity matrix, and $[A_1 | \dots | A_t]$ is a codeword of a Gabidulin code of length bt and minimum rank distance b . Of particular interest are the “unit” subspaces,

$$U_i \triangleq \text{rowsp}[\underbrace{0 \dots 0}_{i-1} | I_b | 0 \dots 0],$$

for all $i \in [M/b]$. Obviously,

$$\sum_{i=1}^{M/b} U_i = \mathbb{F}_q^M.$$

Thus, except for unit subspaces from $U \triangleq \{U_i\}_{i \in [M/b]}$, for every other subspace of the spread, the set U is a recovery set of M/b thick columns.

We are left with the task of finding recovery sets of unit subspaces of the form U_i . For every $i \in [M/b - 1]$, we have

$$U_i \subseteq U_{i+1} + \text{rowsp}[\underbrace{0 \dots 0}_{i-1} | I_b | A | 0 \dots 0],$$

where $A \neq 0$ is a codeword of the above-mentioned Gabidulin code. Finally,

$$U_{M/b} \subseteq U_{M/b-1} + \text{rowsp}[0 \dots 0 | I_b | A],$$

since A is full rank due to the minimum rank distance of the Gabidulin code. Thus, each U_i has a recovery set of size $2 \leq M/b$. ■

The code of Construction B is also a generalization of the simplex code. Indeed, when we take $b = 1$ the resulting generator matrix is that of a simplex code.

Corollary 10. *When $M = 2b$, the code from Construction B is an MDS array code with $r_n = r_s = 2$.*

Proof: The node and symbol locality are trivial since the subspaces associated with thick columns have a pairwise trivial intersection, and therefore the sum of any two such subspaces gives the entire space since $M = 2b$. The code is MDS since it is a $[b \times (q^b + 1), 2b, q^b]$ array code. ■

Up to this point we constructed codes by specifying their generator matrix. We now turn to consider their dual codes by reversing the roles of generator and parity-check matrices. We first require the following simple lemma.

Lemma 11. *Let C be a $[b \times n, M, d]$ array code over \mathbb{F}_q that is full column rank. If the size of the smallest recovery set for a symbol of C is of size ℓ , then the dual code, C^\perp , is a $[b \times n, bn - M, \ell + 1]$ array code. In particular, if the symbol locality of every symbol of C is r_s , then C^\perp is a $[b \times n, bn - M, r_s + 1]$ array code.*

Proof: Let G be a generator matrix for C . The smallest recovery set of size ℓ together with the full column rank property imply that the smallest set of linearly dependent columns of G includes columns from exactly $\ell + 1$ thick columns. Considering G as a parity-check matrix for C^\perp , we obtain that the any non-zero codeword of C^\perp has at least $\ell + 1$ non-zero columns. The rest of the code parameters are trivially obtained. ■

The dual code of the code from Construction A has a small distance $d = 2$, and is therefore not very interesting. However, the code from Construction B presents a more interesting situation.

Lemma 12. *Let C be a code from Construction B. Then its dual, C^\perp , is a $[b \times \frac{[M]}{[1]} / \frac{[b]}{[1]}, b \frac{[M]}{[1]} / \frac{[b]}{[1]} - M, 3]$ array code. Additionally, C^\perp is a perfect array code.*

Proof: The minimum distance follows from Lemma 11 since the locality of all symbols in C is 2. To show that C^\perp is perfect, note that the ball of radius 1 has size

$$\Phi_1 \triangleq 1 + \frac{[M]}{[b]}(q^b - 1) = q^M.$$

Hence,

$$|C^\perp| \cdot \Phi_1 = q^{b \frac{[M]}{[1]} / \frac{[b]}{[1]}},$$

which is equal to the size of the entire space. ■

We note that the code of Lemma 12 has already been described as a perfect byte-correcting code in [6], [14].

At this point we stop to reflect back on Construction A and Construction B. We contend that the two are in fact two extremes of a more general construction using the q -analog of Steiner systems.

Definition 13. *Let F_q be a finite field. A q -analog of a Steiner system (a q -Steiner system for short), denoted $S_q[t, k, n]$, is a set of subspaces, $\mathcal{B} \subseteq \left[\begin{smallmatrix} \mathbb{F}_q^n \\ k \end{smallmatrix} \right]$, such that every subspace from $\left[\begin{smallmatrix} \mathbb{F}_q^n \\ t \end{smallmatrix} \right]$ is contained in exactly one element of \mathcal{B} .*

In light of Definition 13, we note that the subspaces associated with the columns of Construction A form a q -Steiner system $S_q[b, b, M]$. Similarly, the subspaces associated with the columns of Construction B form a q -Steiner system $S_q[1, b, M]$. Both are therefore extreme (and trivial) cases of a more general construction we now describe.

Construction C. *Fix a finite field \mathbb{F}_q , and let $\mathcal{B} \subseteq \left[\begin{smallmatrix} \mathbb{F}_q^M \\ b \end{smallmatrix} \right]$ be a q -Steiner system $S_q[t, b, M]$. Construct an array code whose set of columns are associated with the subspace set \mathcal{B} , each appearing exactly once.*

The main problem with the approach of Construction C is the fact that we need a q -Steiner system. Such systems are extremely hard to find [2], [24], with the only known ones (different $S_2[2, 3, 13]$), found by computer search [2].

An alternative approach uses a structure that is “almost” a q -Steiner system, and is more readily available – a subspace transversal design (see [8]).

Definition 14. *Let \mathbb{F}_q be a finite field. A subspace transversal design of group size $q^m = q^{n-k}$, block dimension k , and strength t , denoted by $\text{STD}_q(t, k, m)$ is a triple $(\mathcal{V}, \mathcal{G}, \mathcal{B})$, where*

- 1) $\mathcal{V} \triangleq \left[\begin{smallmatrix} \mathbb{F}_q^n \\ 1 \end{smallmatrix} \right] \setminus \mathcal{V}_0^{(n,k)}$, called the points, where $\mathcal{V}_0^{(n,k)}$ is defined to be the set of all 1-dimensional subspaces of \mathbb{F}_q^n all of whose vectors start with k zeros, and where $|\mathcal{V}| = \left[\begin{smallmatrix} k \\ 1 \end{smallmatrix} \right] q^m$.
- 2) \mathcal{G} is a partition of \mathcal{V} into $\left[\begin{smallmatrix} k \\ 1 \end{smallmatrix} \right]$ classes of size q^m , called the groups.

- 3) $\mathcal{B} \subseteq \left[\begin{smallmatrix} \mathbb{F}_q^n \\ k \end{smallmatrix} \right]$, called the blocks, is a collection of subspaces that contain only points from \mathcal{V} , with $|\mathcal{B}| = q^{mt}$.
- 4) Each block meets each group in exactly one point.
- 5) Each t -dimensional subspace of \mathbb{F}_q^n , with points only from \mathcal{V} , which meets each group in at most one point, is contained in exactly one block.

An $\text{STD}_q(t, k, m) = (\mathcal{V}, \mathcal{G}, \mathcal{B})$ is called resolvable if the set \mathcal{B} may be partitioned into sets $\mathcal{B}_1, \dots, \mathcal{B}_s$, called parallel classes, where each point is contained in exactly one block of each parallel class \mathcal{B}_i .

Unlike q -Steiner systems, subspace transversal designs are known to exist in a wide range of parameters, as shown in the following theorem [8].

Theorem 15. [8, Th. 7] *For any $1 \leq t \leq k \leq m$, and any finite field \mathbb{F}_q , there exists a resolvable $\text{STD}_q(t, k, m) = (\mathcal{V}, \mathcal{G}, \mathcal{B})$, where the block set \mathcal{B} may be partitioned into $q^{m(t-1)}$ parallel classes, each one of size q^m , such that each point is contained in exactly one block of each parallel class.*

Construction D. *Fix a finite field \mathbb{F}_q , $M \geq 2b$, and let $(\mathcal{V}, \mathcal{G}, \mathcal{B})$ be a $\text{STD}_q(t, b, M - b)$ with parallel classes $\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_s$. Construct the following two array codes:*

- An array code C_{par} whose set of columns are associated with the subspaces in a single parallel class, \mathcal{B}_i , each appearing exactly once.
- An array code C whose set of columns are associated with the subspaces in \mathcal{B} , each appearing exactly once.

The code C_{par} is in fact an auxiliary code we shall use to prove the parameters of the code C , and is perhaps of interest on its own.

Theorem 16. *Let C_{par} be the code from Construction D. Then C_{par} is a $[b \times q^{M-b}, M, q^{M-b} - q^{M-2b}]$ array code, with $2^b - 1$ codewords of full weight q^{M-b} , and the other non-zero codewords of weight $q^{M-b} - q^{M-2b}$. Moreover, the symbol locality of C_{par} is $r_s = 2$, and its node locality is*

$$r_n = \begin{cases} 3 & q = 2, \\ 2 & q > 2. \end{cases}$$

Proof: The size and dimension of the array code follow from Theorem 15. The rest of the proof follows the same logic as the proof of Theorem 8.

Denote $u \triangleq q^{M-b}$. Consider an $M \times bu$ generator matrix G for C_{par} . It contains u thick columns, each made up of b columns. Let G_i , $i \in [u]$, be the $M \times b$ submatrix of G containing the b columns of the i th thick column, i.e., $G = (G_1 | G_2 | \dots | G_u)$.

We now take each G_i , $i \in [u]$, and construct from it an $M \times (q^b - 1)$ matrix we call G_i^{ext} , whose columns are the column space of G_i except for $\vec{0}$. We concatenate those to obtain the $M \times u(q^b - 1)$ matrix

$$G^{\text{ext}} \triangleq (G_1^{\text{ext}} | G_2^{\text{ext}} | \dots | G_u^{\text{ext}}).$$

Since we used a single parallel class, the columns of G^{ext} contain each possible vector exactly once, except for columns beginning with b zeros. In other words, the

subspaces of dimension b that correspond to the thick columns of G , together with the subspace of dimension $M - b$ of all vectors starting with b zeros, form a partition of the non-zero vectors of \mathbb{F}_q^M .

We now observe that a row of G_i^{ext} is $\bar{0}^T$ iff it is $\bar{0}^T$ in G_i . Additionally, a non-zero row of G_i^{ext} contains exactly q^{b-1} occurrences of each non-zero element of \mathbb{F}_q . It is now a matter of simple counting, to obtain that each of the first b rows of G^{ext} has all of its $u = q^{M-b}$ thick columns non-zero, and the remaining lower $M - b$ rows of G^{ext} have exactly $q^{M-b} - q^{M-2b}$ non-zero thick columns in each row.

Finally, consider a linear combination of the rows of G that involves rows i_1, i_2, \dots, i_ℓ , all with non-zero coefficients, and resulting in a row \bar{v}^T . As in the proof of Theorem 8, let us replace row i_ℓ of G with \bar{v}^T to obtain a new generator matrix G' . Again, the subspaces the correspond to the thick columns of G' induce a partition of the non-zero vectors of \mathbb{F}_q^M into subspaces of dimension b and a single subspace of dimension $M - b$. Therefore, we conclude that the resulting row corresponds to an array codeword of weight either q^{M-b} or $q^{M-b} - q^{M-2b}$ depending on whether $i_1, \dots, i_\ell \in [b]$ or not. This gives us a total of $q^b - 1$ codewords in C_{par} of weight q^{M-b} , and the remaining non-zero codewords of weight $q^{M-b} - q^{M-2b}$.

To complete the proof, the symbol locality is $r_s = 2$, since any column of G may be easily be given as a sum of two other columns of G (which must also reside in distinct thick columns), due to the partition of \mathbb{F}_q^M discussed above. To prove the node locality we recall that any thick column of G corresponds to a lifted MRD codeword, i.e., $(I_b|A)^T$, where A is a codeword of a linear MRD code of dimension $M - b$. When $q = 2$, we can recover $(I_b|A)^T$ by noting that

$$(I_b|A)^T = (I_b|A')^T + (I_b|A + A')^T + (I_b|0)^T,$$

where A' is a codeword of the lifted MRD code, $A' \neq A$, and where we use the fact that $M - b \geq 2$. When $q > 2$, let $\alpha \in \mathbb{F}_q$, $\alpha \neq 0, 1$. Then we can recover $(I_b|A)^T$ by noting that

$$(I_b|A)^T = \alpha^{-1}(I_b|\alpha A) + (\alpha - 1)\alpha^{-1}(I_b|0)^T,$$

thus proving $r_n = 2$ for $q > 2$. ■

Corollary 17. When $M = 2b$, the code C_{par} from Construction D is an MDS array code with $r_n = r_s = 2$.

Proof: The node and symbol locality are trivial since the subspaces associated with thick columns have a pairwise trivial intersection, and therefore the sum of any two such subspaces gives the entire space since $M = 2b$. The code is MDS since it is a $[b \times q^b, 2b, q^b - 1]$ array code. ■

Corollary 18. Let C_{par} be the code from Construction D. Then its dual code, C_{par}^\perp is a $[b \times q^{M-b}, bq^{M-b} - M, 3]$ array code that is asymptotically perfect.

Proof: The parameters of the code follow from Lemma 11 and from the proof of Theorem 16. Note that the size of a ball of radius 1 is equal to

$$\Phi_1 \triangleq 1 + q^{M-b}(q^b - 1).$$

The size of the entire space is $q^{bq^{M-b}}$. Then

$$\begin{aligned} \frac{|C_{\text{par}}^\perp| \cdot |\Phi_1|}{q^{bq^{M-b}}} &= \frac{q^{bq^{M-b}-M}(1 + q^{M-b}(q^b - 1))}{q^{bq^{M-b}}} \\ &= \frac{1 + q^M - q^{M-b}}{q^M} = 1 + q^{-M} - q^{-b}, \end{aligned}$$

and this ratio tends to 1 when $b, M \rightarrow \infty$, implying the code family is asymptotically perfect. ■

Example 19. Let $b = 3$, $M = 6$, $q = 2$. A generator matrix G for the $[3 \times 8, 6, 7]$ MDS array code C_{par} from Construction D is given by

$$G = \left(\begin{array}{c|c|c|c|c|c|c|c} 100 & 100 & 100 & 100 & 100 & 100 & 100 & 100 \\ 010 & 010 & 010 & 010 & 010 & 010 & 010 & 010 \\ 001 & 001 & 001 & 001 & 001 & 001 & 001 & 001 \\ 000 & 100 & 001 & 010 & 101 & 011 & 111 & 110 \\ 000 & 010 & 101 & 011 & 111 & 110 & 100 & 001 \\ 000 & 001 & 010 & 101 & 011 & 111 & 110 & 100 \end{array} \right).$$

□

We now move on to examine the second code of Construction D. To avoid degenerate cases, we consider only $t \geq 2$.

Theorem 20. Let C be the code from Construction D, with $t \geq 2$. Then C is a $[b \times q^{(M-b)t}, M, d]$ array code

$$d = q^{(M-b)(t-1)}(q^{M-b} - q^{M-2b}).$$

The symbol and node locality of the code satisfy $r_s = 1$, and $r_n \geq 2$. Its symbol availability is $t_s = q^{(M-b)(t-1)} - 1$.

Proof: The codeword size, as well as the minimum distance follow immediately by noting that there are $q^{(M-b)(t-1)}$ parallel classes, and a generator matrix for C is simply the concatenation of generators for C_{par} (for each of the parallel classes). The minimum distance then follows from Theorem 16.

Additionally, each point (i.e., a column of G) is contained exactly once in each of the $q^{(M-b)(t-1)}$ parallel classes in a single subspace (i.e., the column span of a thick column of G). Thus, as long as $t \geq 2$, the symbol locality is $r_s = 1$, and the availability is $t_s = q^{(M-b)(t-1)} - 1$. Trivially, by the properties of the subspace transversal design, no subspace associated with a thick column appears twice, and hence $r_n \geq 2$. ■

REFERENCES

- [1] R. Bhagwan, K. Tati, Y. Cheng, S. Savage, and G. M. Voelker, "Total recall: system support for automated availability management," *Networked Sys. Design and Implem. (NSDI)*, pp. 337–350, 2004.
- [2] M. Braun, T. Etzion, P. R. J. Östergård, A. Vardy, and A. Wassermann, "Existence of q -analogs of Steiner systems," *Forum of Mathematics, Pi*, vol. 4, no. e7, pp. 1–14, 2016.
- [3] A. Datta and F. Oggier, "An overview of codes tailor-made for networked distributed data storage," *arXiv:1109.2317*, Sep. 2011.
- [4] A. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Trans. Inform. Theory*, vol. 56, no. 9, pp. 4539–4551, Sep. 2010.
- [5] A. Dimakis, K. Ramchandran, Y. Wu, and C. Suh, "A survey on network codes for distributed storage," *Proc. of the IEEE*, vol. 99, pp. 476–489, 2011.

- [6] T. Etzion, "Perfect byte-correcting codes," *IEEE Trans. Inform. Theory*, vol. 44, no. 7, pp. 3140–3146, Nov. 1998.
- [7] T. Etzion and N. Silberstein, "Error-correcting codes in projective space via rank-metric codes and Ferrers diagrams," *IEEE Trans. Inform. Theory*, vol. 55, no. 7, pp. 2909–2919, Jul. 2009.
- [8] —, "Codes and designs related to lifted MRD codes," *IEEE Trans. Inform. Theory*, vol. 59, no. 2, Feb. 2013.
- [9] T. Etzion and A. Vardy, "Error-correcting codes in projective space," *IEEE Trans. Inform. Theory*, vol. 57, no. 2, pp. 1165–1173, Feb. 2011.
- [10] E. M. Gabidulin and N. Pilipchuk, "Multicomponent network coding," in *WCC 2011-Workshop on coding and cryptography, Paris, France*, Apr. 2011, pp. 443–452.
- [11] S. Ghemawat, H. Gobioff, and S.-T. Leung, "The Google file system," in *ACM SIGOPS operating systems review*, vol. 37, no. 5, 2003, pp. 29–43.
- [12] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Trans. Inform. Theory*, vol. 58, no. 11, pp. 6925–6934, Nov. 2012.
- [13] H. D. L. Hollmann, "Storage codes; coding rate and repair locality," in *Proceedings of the Int. Conf. on Computing, Networking and Communications (ICNC), San Diego, CA, USA*, Jan. 2013, pp. 830–834.
- [14] S. J. Hong and A. M. Patel, "A general class of maximal codes for computer applications," *IEEE Trans. Comput.*, vol. C-21, pp. 1322–1331, 1972.
- [15] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li, and S. Yekhanin, "Erasure coding in Windows Azure storage," in *Proc. USENIX ATC 12, Boston, MA, USA*, 2012, pp. 15–26.
- [16] G. Kamath, N. Prakash, V. Lalitha, and P. Kumar, "Codes with local regeneration and erasure correction," *IEEE Trans. Inform. Theory*, vol. 60, no. 8, pp. 4637–4660, Aug. 2014.
- [17] G. Kamath, N. Silberstein, N. Prakash, A. Rawat, V. Lalitha, O. Koyluoglu, P. Kumar, and S. Vishwanath, "Explicit MBR all-symbol locality codes," in *Proceedings of the 2013 IEEE International Symposium on Information Theory (ISIT2013), Istanbul, Turkey*, Jul. 2013, pp. 504–508.
- [18] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*. North-Holland, 1978.
- [19] F. Oggier and A. Datta, "Self-repairing homomorphic codes for distributed storage systems," in *Proc. INFOCOM, Shanghai, China*, Apr. 2011, pp. 1215–1223.
- [20] N. Raviv and T. Etzion, "Distributed storage systems based on intersecting subspace codes," in *Proceedings of the 2015 IEEE International Symposium on Information Theory (ISIT2015), Hong Kong, SAR China*, Jun. 2015, pp. 1462–1466.
- [21] A. Rawat, O. Koyluoglu, N. Silberstein, and S. Vishwanath, "Optimal locally repairable and secure codes for distributed storage systems," *IEEE Trans. Inform. Theory*, vol. 60, no. 1, pp. 212–236, Jan. 2014.
- [22] A. Rawat, D. Papailiopoulos, A. Dimakis, and S. Vishwanath, "Locality and availability in distributed storage," in *Proceedings of the 2014 IEEE International Symposium on Information Theory (ISIT2014), Honolulu, HI, USA*, Jun. 2014, pp. 631–635.
- [23] S. C. Rhea, P. R. Eaton, D. Geels, H. Weatherspoon, B. Y. Zhao, and J. Kubiatowicz, "Pond: The oceanstore prototype," in *Proc. 2th USENIX Conference on File and Storage Technologies (FAST), San Francisco, CA, USA*, vol. 3, Mar. 2003, pp. 1–14.
- [24] M. Schwartz and T. Etzion, "Codes and anticodes in the Grassman graph," *J. Combin. Theory Ser. A*, vol. 97, no. 1, pp. 27–42, Jan. 2002.
- [25] I. Tamo, Z. Wang, and J. Bruck, "Zigzag codes: MDS array codes with optimal rebuilding," *IEEE Trans. Inform. Theory*, vol. 59, no. 3, pp. 1597–1616, Mar. 2013.
- [26] J. H. van Lint and R. M. Wilson, *A Course in Combinatorics, 2nd Edition*. Cambridge Univ. Press, 2001.